THE SECRET WORLD OF OSCILLOSCOPE PROBES

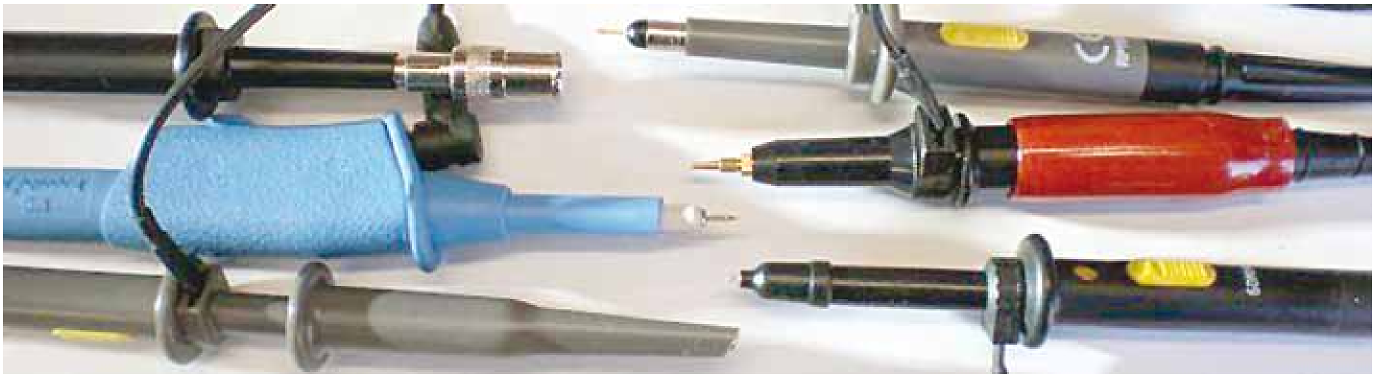

Ever wondered how scope probes really work? Most textbooks treat scope probes as a combination of a resistive divider in combination with capacitors to provide an extended frequency response. But as will be revealed, the reality is that they are much more complex in principle. Read on.

By Doug Ford

The oscilloscope is an essential tool for anyone working in electronics. Whether you're working in electronics service, production, testing R&D or in your home workshop, you need an oscilloscope.

If you listen to a bunch of technical people chatting about their scopes, they'll talk about their bandwidth or whether they have colour displays, depth of memory or portability but the probes rarely get mentioned.

In fact, most users don't think about their probes until they hear the sickening crunch underfoot which tells them they shouldn't have left them dangling off the bench onto the floor.

There are many varieties of "specialist" probes: active-FET probes, differential-floating probes, current-sensing probes are just some we could mention.

They all have their uses but by far the most common is the "times ten" (x10) passive voltage probe. Typically, you're given two of them free with every oscilloscope.

But how much do you really know about these probes?

A few hours of Googling will yield countless explanations about basic operation (voltage division and capacitance compensation) but you are unlikely to find explanations which show the probe's transmission-line properties.

Nor will you find any adequate description of the design differences between inexpensive 40MHz probes and much dearer 350MHz probes.

Conventional explanations

Conventional wisdom explains the operation of a x10 probe with the equivalent circuit in Fig.1 (above right).

The scope's input impedance is assumed to be 1MΩ in parallel with a small capacitance (somewhere between 10pF and 50pF). Low-bandwidth scopes generally have higher input capacitances.

The capacitance of the probe cable may be from 60pF (for a high bandwidth probe) to 200pF (for a pretty average probe).

The factor-of-ten voltage division is determined at lower frequencies by the divider formed by the 9MΩ resistor in the tip of the probe and the 1MΩ scope input resistance.

The compensation capacitor across the 9MΩ probe resistor is trimmed to be 1/9th the combined capacitances of the scope input and the probe cable. In the case above, the scope-plus-cable

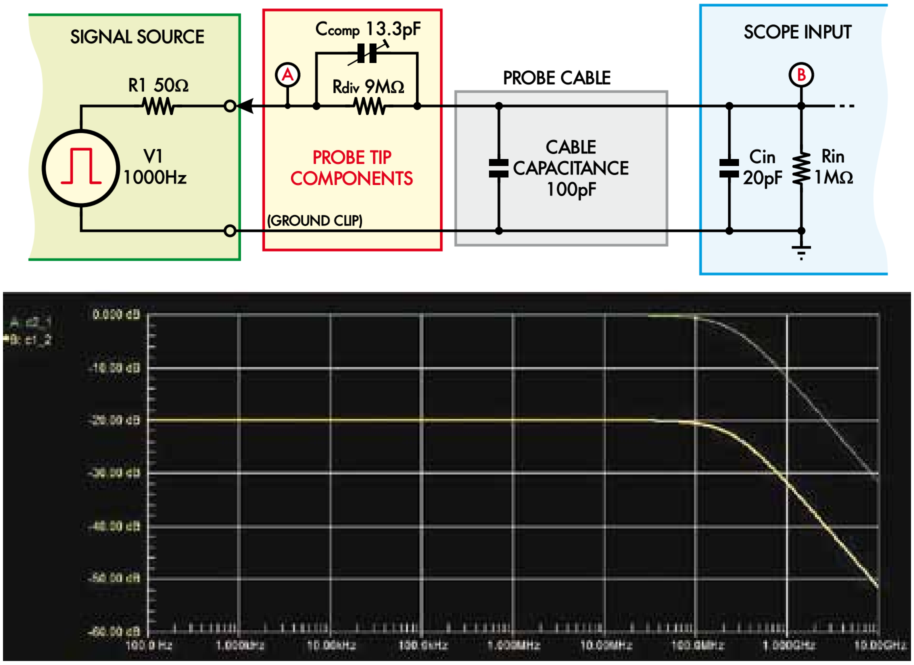

Fig.1: Circuit and response of x10 probe ("Conventional" explanation).

capacitance is 120pF, so the compensation cap is trimmed to (120/9) = 13.3pF.

When the capacitive divider formed by C_{COMP} and C_{IN}//C~~CBL~~ has the same 1:10 ratio as the resistive divider formed by Rin and Rdiv, the frequency response of the probe should be flat from DC to ultraviolet.

The only limitation to high-frequency bandwidth should be the interaction of the source impedance (shown here as 50Ω) with the effective capacitance of the probe tip (12pF), giving a –3dB point of 265MHz.

Note that the frequency scale of the simulation extends from 10Hz to 10GHz. We don't want to miss any interesting artefacts, do we?

Trimming the compensation capacitor

The effect of trimming the compensation capacitor on frequency response is shown in Fig.2 The capacitor has been varied from 8pF to 18pF in 1pF steps.

Note that the gain is unaffected at frequencies below 300Hz but gain errors in the 3kHz ~ 100MHz range are large and consistent.

Oscilloscopes are fitted with an internal square-wave generator which feeds a "calibration" terminal on the front panel.

This calibration signal is provided specifically for the purpose of trimming probes. The calibration signal frequency is usually 1kHz with an amplitude of 1V peak-to-peak. The probe is connected to the calibration terminal and adjusted to achieve the "squarest" waveform display.

Anyone who has trimmed a x10 probe will be familiar with the scope waveform seen during trimming, as in Fig.3.

While Fig.1 shows the compensation trim capacitor connected across the 9MΩ probe resistor this is actually very rare.

More typically, the capacitor across the 9MΩ resistor has a fixed value and trimming is achieved by a trimmer connected in parallel with the probe cable and scope input capacitances, as shown in Fig.4.

At this stage, there doesn't appear to be much difference between probes with tip-end or scope-end trimming. Both types of probe are available, with bandwidths from 20MHz to 300MHz.

However, higher bandwidth probes

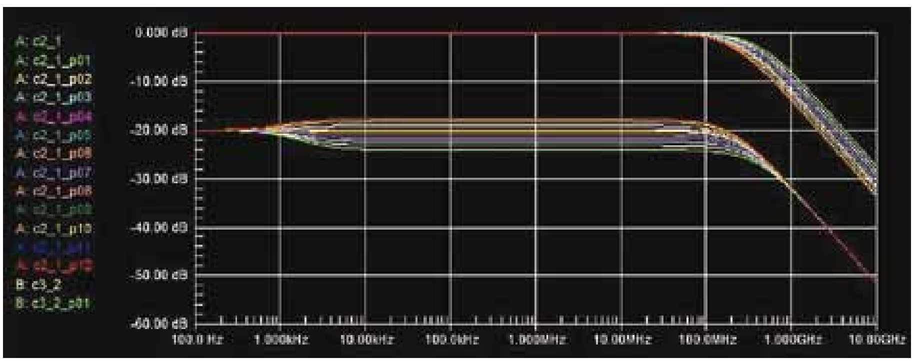

Fig.2: How compensation trimming affects frequency response.

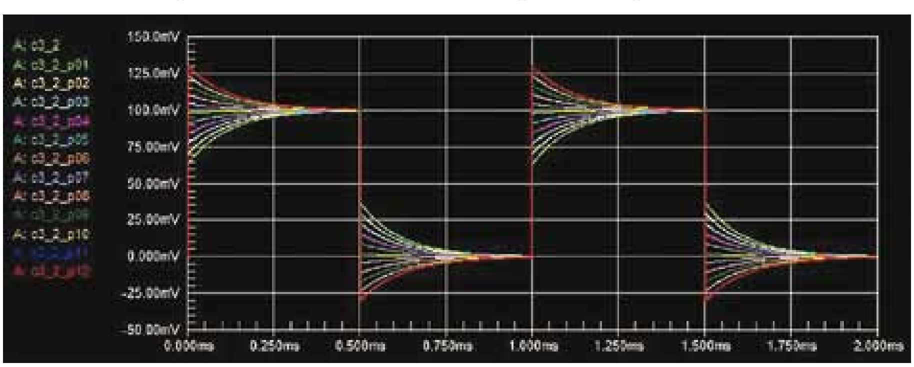

Fig.3: Waveforms seen during compensation trimming of a 1kHz square-wave.

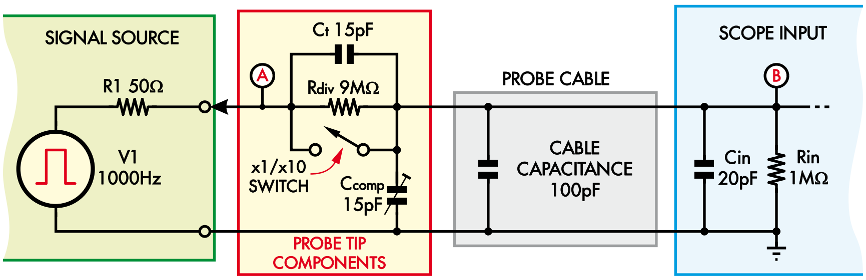

Fig. 4: Probe circuit with fixed tip capacitor.

Fig.5: compensation trimmer at the scope end (left) and probe end (right).

(350MHz and higher) tend to have their compensation trimmers at the scope end of the cable.

So far, we have given a fairly simple description of probe operation using standard textbook explanations. But this ignores the fact that the probe's cable is NOT a simple lumped capacitance; it is a transmission line!

The probe's coaxial cable has length, distributed inductance and capacitance, propagation delay and signal reflections from unterminated ends. What's the effect of these properties on the behaviour of a probe?

So let's replace the lumped cable capacitance in our previous simulation with a transmission line and see what happens!

Simulator software

CircuitMaker is a schematic layout and simulator program originally released by Microcode.

I've been told that Microcode bought the Autotrax franchise from Protel in the early 90's. In 1998, Protel bought Microcode, then changed its name to Altium in 2001. So, CircuitMaker became an Altium product, until Altium discontinued it in 2001.

This is a pity, because this excellent simulator was bundled with the PCB drafting program Traxmaker (a Windows version of Autotrax) and a Gerber file reader at a very reasonable price.

Simulating a probe's cable

We'll replace the single 100pF cable capacitance with a transmission line in the circuit simulator. The circuit simulator can simulate any transmission line, but we need to make a few guesses about the circuit values to enter into the simulator.

Typical probe cables are around 1.2m long, although they can be up to 1.8m. The total capacitance of my 250MHz probes is 85pF, according to their manufacturer's specifications. The specified capacitance is 128pF for my 60MHz probes, although these actually measured closer to 170pF. We will use 100pF in simulations for now, to maintain parity with the previous simulations of Fig.1 and Fig.4. Our cable capacitance will thus be 83pF/m for a 1.2m cable.

We will assume that the cable's characteristic impedance is 50Ω for the moment. The cable inductance (per unit length) can be calculated from: $Z_O = \sqrt{(L/C)}$, where L = inductance per unit length and C = capacitance per unit length. It doesn't matter what your unit length is; we're using metres here. The calculated inductance, for a 50Ω line with 83pF/m capacitance, is 208nH/m.

Since we haven't changed the 100pF cable capacitance, we don't need to change the 15pF tip capacitor or 15pF compensation trim.

These values were punched into the transmission line shown in Fig.6:

The result is an awful frequency response! The effects of reflections from the unterminated transmission line will give huge response variations at the scope above 20MHz (green trace). The effects of probe loading on the signal generator (yellow trace) are similarly large.

So what do probe and scope designers do to address this problem?

I have looked inside several scopes and probes over the years. And I've trodden on a few probes in my time, resulting in some sad post-mortems and furtive probe replacements.

Most probes have a discrete low-value resistor built into the probe tip extremity, located at the tip in front of the 9MΩ divider resistor and x1/x10 switch.

I measured the end-to-end resistance of some probes (in x1 setting) and found values in the range 180Ω ~ 270Ω. OK, we will include some probe-tip resistance, say 250Ω in the simulation.

Similarly, I have seen that in some older scopes, there is a series 50Ω resistor between the BNC input socket and the range switch. We will include this, as well. See Fig.7.

The frequency response (green) is obviously smoother than in Fig.4 and the loading effect on the source (yellow) is

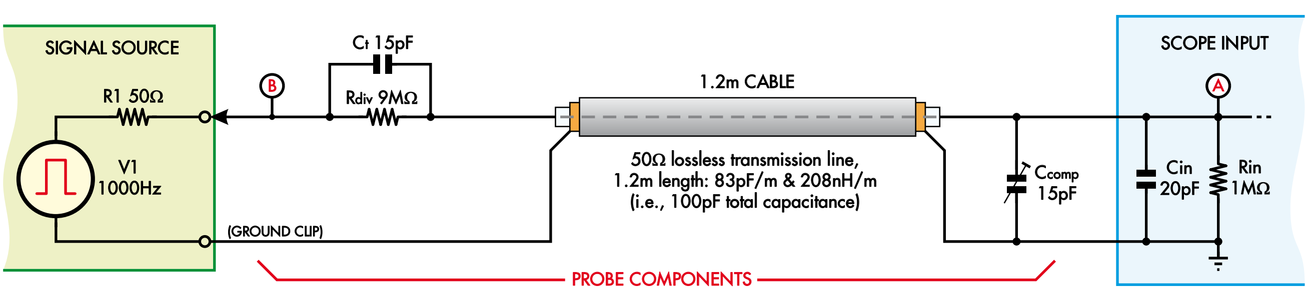

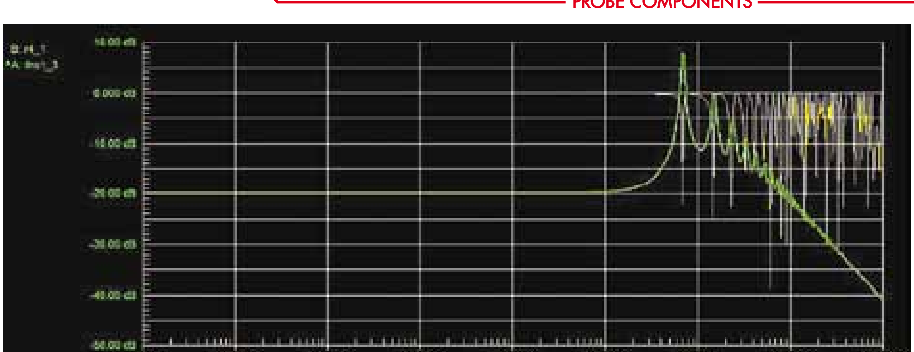

Fig.6: Simple transmission-line model: Circuit diagram and frequency response.

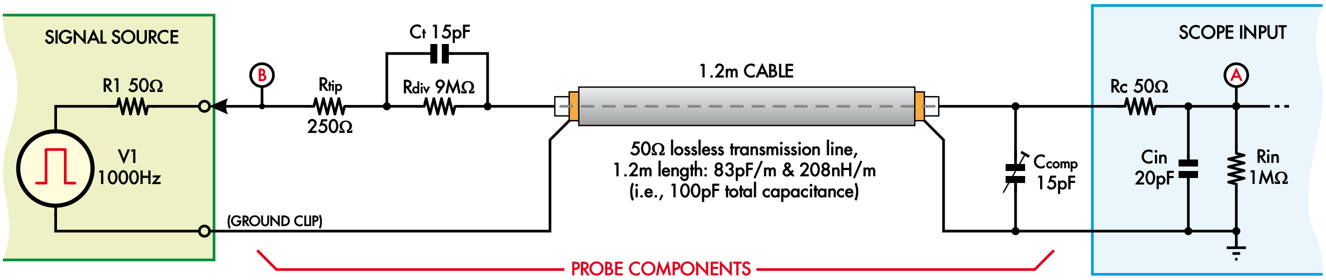

Fig.7: Simple TL model with added probe & scope resistances: circuit and frequency response

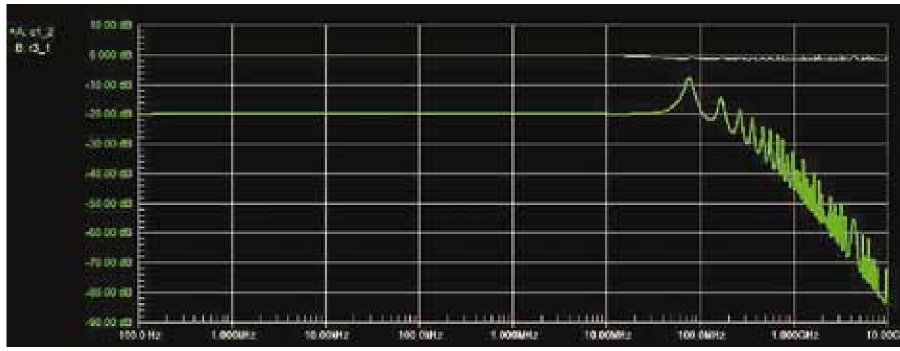

lower. But the usable bandwidth is still less than 40MHz. Even if there was a clever way to smooth the response, it would still only get to maybe 100MHz before rolling off.

Tweaking the compensation capacitor has little effect on the frequency response or the transmission-line resonance effects. So it is obvious that the transmission-line characteristics of the probe cable are potentially responsible for some serious bandwidth and frequency-response limitations. So, what is the secret behind the design of my 250MHz probes, and even my junkbox 60MHz probes? How DO probe manufacturers manage to get extreme bandwidths from probes?

I tried all kinds of tricks in simulator-land to see how the transmission line could be tamed and how the response could be extended. I tried variations to the cable's characteristic impedance, various component combinations at the tip or at the scope end of the cable; all to little effect.

I eventually resorted to examination of the cable from a defunct probe. I discovered that the cable centre core had a surprisingly high resistance.

I dissected the cable further and was surprised to discover that the core wire appeared to be very thin resistance wire, with a resistance of around 100 ~ 200Ω per metre! See Fig.8.

This very fine core wire appears to be made from a single strand and is "crinkled" – presumably to provide tolerance to repeated flexing. I'm guessing that the white foam core insulation gives low dielectric loss, while the black PVC around the foam gives mechanical support to the foam (and no, the black stuff isn't conductive. I checked!).

The high resistance of the core wire was the clue I needed. This coax cable is NOT low-loss; it has been made deliberately lossy, to reduce the effects of end-to-end transmission-line reflections!

I now wanted to know the identity of the unknown, unsung genius who developed this trick.

So, back to simulator-land. This time, we'll give the coaxial cable a resistance of 165Ω/m (200Ω total). We'll also reduce the value of the probe-tip resistor from 250Ω to 50Ω. The overall probe series resistance is still 250Ω, as before.

Also, I'm pretty sure that most modern scopes don't use 50Ω series resistors any more, because modern high bandwidth scopes have very low input capacitances (10pF ~ 15pF). This renders the scope's 50Ω series terminator pretty useless at frequencies around 80MHz, where transmission-line end-to-end resonance is most problematic. It's irrelevant so let's get rid of it from simulations.

Fig.9 shows the magic result: a smooth and monotonic response, which is -3dB down at 65MHz with no nasty reflections or response anomalies – just a smooth, usable response!

Even more interesting: the response of this simulation conforms quite nicely to the behaviour of a typical 60MHz probe!

OK, so now we know the secret to designing a probe: use lossy transmission-line cable! But how can the response be extended?

First: I'll assume that modern high-bandwidth scopes don't have 50Ω series termination.

Secondly: I'll use the manufacturers' specs for a 100MHz oscilloscope and 250MHz probe in the simulator.

Thirdly: I'll assume a low-impedance source, instead of the 50Ω source impedance used so far.

Fourthly: when I dissected the scope-end compensation trim of the cable shown in Fig.8, I found that the trimmer capacitor was connected in series

Fig.8: probe cable dissection – note the crinkled inner wire.

Transmission Lines

Transmission lines may take many physical forms: They be in the form of single conductors near a ground return, such as copper tracks on PC boards, PC striplines and single-wire rural phone lines. They may be in the form of wire pairs, such as figure-8 cable, twisted wire pairs or overhead power transmission lines. They may be in the form of coaxial cables, whether single-conductor, stranded conductor or shielded twisted pairs.

As a rough rule of thumb, wire conductors will begin to exhibit transmission-line effects when their length becomes greater than one-tenth of a wavelength while conductors longer than a quarter wavelength show definite transmission-line effects.

Mains power lines operating at 50Hz are treated as transmission lines if their length exceeds a few hundred kilometres. Phone lines with 3kHz bandwidth are treated as transmission lines if they are longer than a few kilometres. At 10MHz, any conductor longer than 30cm must be treated as a transmission line!

A property of a transmission line is its characteristic impedance. When a transmission line is loaded at its far end by a resistor of the same value as its characteristic impedance, all signals fed into the line are absorbed by this resistor. If the load at the far end is not the same as the line's characteristic impedance, signals will be reflected from the far end back to the signal source.

If the line is fed with signals via a resistance equal to the characteristic impedance, it doesn't matter if the far end is not terminated by the correct resistance; Any reflections from the far end will be absorbed by the source resistance.

If a line is terminated by mismatched impedances at both ends (for example, driven at one end from a very low impedance source,and open-circuit or short-circuit at the far end) then signals can ping-pong up and down the line many times before they are slowly absorbed by line losses.

Video distribution systems, which send high-frequency signals through long coaxial cables, terminate both ends of each cable. Signals are sent into a cable via a series terminating resistor and the far end of the cable is terminated by a resistor in the appliance (TV or whatever). This system ensures that a cable is terminated even when an appliance is unplugged from the far end.

The transmission-line characteristics (including characteristic impedance) of a conductor are defined by four basic properties of the wire:

- R, the resistance per unit length (Ω/m)
- L, the inductance per unit length (H/m)
- G, the conductance of the dielectric (insulation) per unit length (m/Ω)
- C, the capacitance per unit length (F/m)

The conductor resistance (R) and insulation conductance (G) determine the losses in the transmission line. The conductance is usually low, but can become very significant in coaxial cable if the insulation becomes waterlogged.

You can calculate the line's characteristic impedance (Z_0) from the inductance and capacitance: $Z_0 = \sqrt{(L/C)}$.

In a coaxial cable, L and C are defined by the cable geometry and materials (core diameter, outer diameter and insulation material). If you make the core wire smaller, you increase its inductance and reduce its capacitance, so the characteristic impedance becomes higher. This is why 75Ω coax has a smaller wire diameter than 50Ω coax of similar size.

Low-loss coaxial cables usually use foamed insulation around the core, rather than solid insulation. The gas in the foam reduces the insulation's dielectric constant, reducing capacitance. This allows thicker core wire to be used to achieve the right characteristic impedance, giving lower resistance and lower loss. In addition, the foam insulation can sometimes have lower conductivity (lower loss) than its solid counterpart – at least, until moisture seeps in…

When you are simulating or measuring the effects of a cable at low frequencies where no transmission-line effects are seen, you will be dealing with the "bulk" cable properties. The bulk properties of total resistance, total capacitance and total inductance will be all you require to determine cable effects.

For example, if you're feeding audio signals into a 100m cable (with 100pF/m capacitance) from a 100Ω output source, you would estimate that the high frequency response would be 3dB down at F = 1/2πRC (where R=100Ω and C = 10nF), or around 160kHz; good enough for audio!

However, if you were actually going to feed a 160kHz signal down this same cable, you might want to see if transmission-line effects are likely. Calculate the wavelength from: Wavelength = Velocity/Frequency. Velocity of signals in a cable are around 80% of light-speed (rule of thumb!) or about 250 million meters per second. You don't need much precision for such calculations; just enough information to tell you if you DO need to resort to more elaborate analysis!

At 160kHz, one wavelength = 250,000,000/160,000 = 1500-odd metres. So your 100m cable is one-fifteenth of a wavelength long; You might not have to treat it as a transmission line at 160kHz but you certainly would if your signal had higher harmonics which needed to be preserved.

Transmission-line effects can be simulated and/or calculated by dividing the line into many smaller segments. The inductance, resistance, conductance and capacitance of each segment is given by "quantity per unit length" times cable length, divided by the number of segments. This approach is called the "lumped parameter" method. The equivalent circuit of a lumped-parameter transmission line is shown below.

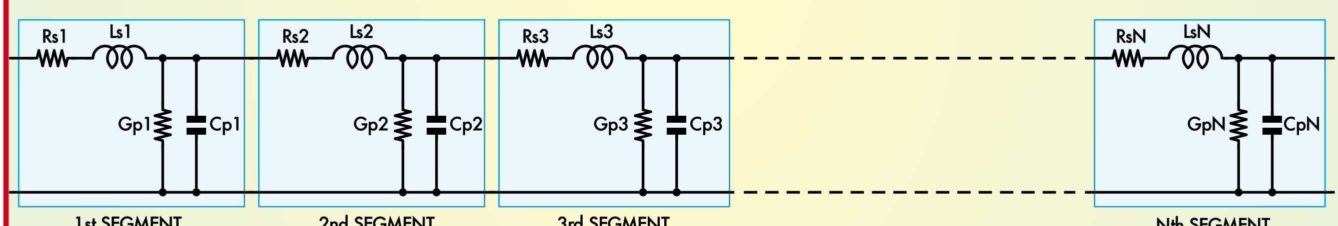

The number of segments (lumps?) you use for your simulation will determine how closely it corresponds with reality. Ten segments will give only moderate accuracy; Several hundred segments will give a very high degree of accuracy to simulations and calculations, but netlist size and computation time can become prohibitive.

Transmission lines are generally modelled in SPICE simulators by matrix mathematics and recursive convolution, rather than by the lumped-parameter approach. These methods require much less computation time than lumped-parameter methods. The maths is beyond me, but the transmission-line model used by CircuitMaker certainly responds correctly to "test questions" which I've posed.

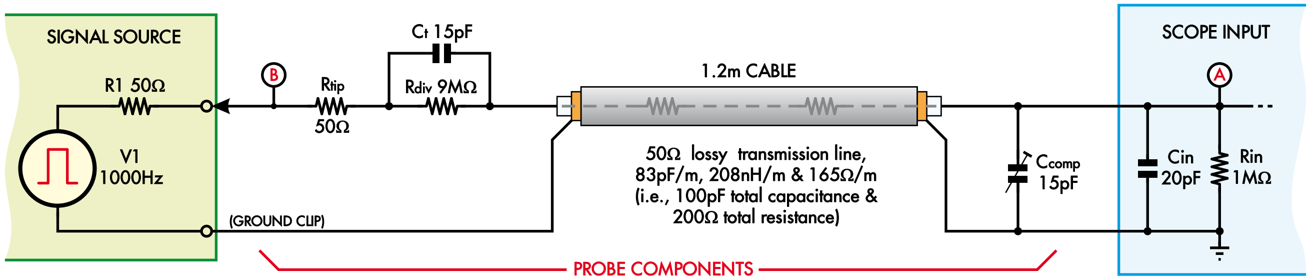

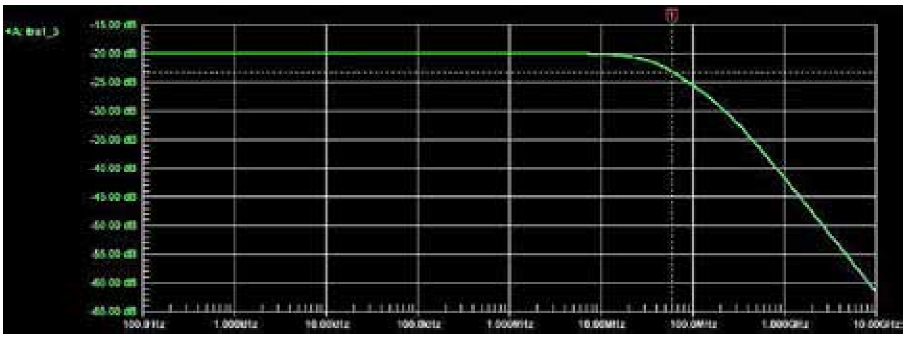

Fig.9: "lossy transmission-line" model and frequency response

with a 68Ω resistor. I'll include this resistor in simulations and find out what it does.

Then we'll juggle the series resistance of the transmission line in the simulator to see what happens!

A 100MHz scope has an input capacitance of 15pF, so we'll use this value at C_{IN}.

My 250MHz probe has a specified capacitance of 85pF (x1 setting), so we'll set the transmission line capacitance to 71pF/m. For a 50Ω cable, the calculated inductance must be 177nH/m.

This probe has a specified capacitance of 15pF (x10 setting), so we'll leave the value of the tip capacitor at 15pF. However, we'll increase the value of the compensation trim to 35pF, because of the lower cable capacitance (85pF vs 100pF).

The simulator circuit using these values is shown in Fig.10.

The transmission-line resistance in this circuit was varied from 50Ω/m to 200Ω/m. This is the kind of experiment where simulators become so incredibly useful. It would be a horribly expensive exercise to obtain the various lossy cables which would be needed to conduct this series of experiments at the test bench.

The effect of varying the cable resistance over the range 50Ω ~ 200Ω per meter can also be seen in Fig.10. Low resistances (50Ω/m) allow transmission-line reflections to build up, giving a peaking response. Larger resistances (200Ω/m) give an overdamped, sagging response.

The optimum cable resistance was found to be around 115Ω/m. This gave a response which is substantially flat to nearly 600MHz!

The real bandwidth of my 250MHz probes would be 250MHz, rather than the 600MHz shown by the simulator. I haven't simulated the small stray capacitances from each component to ground or the stray capacitance across each component, which would reduce the real bandwidth.

The resistor in series with the

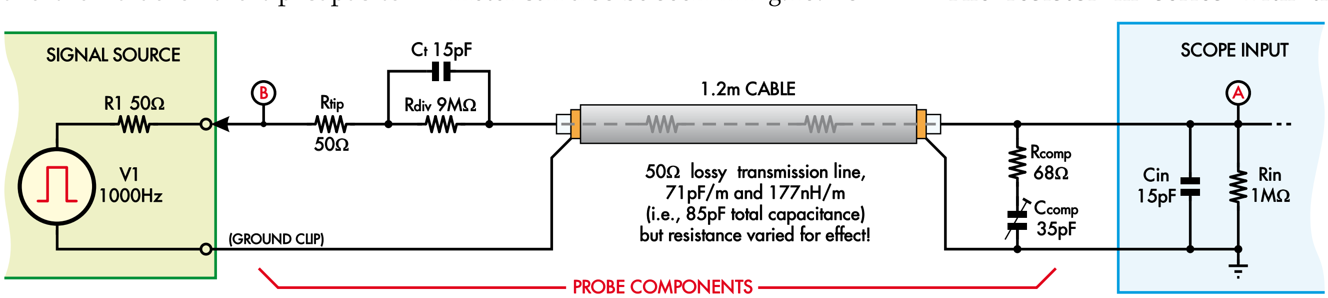

Fig.10: circuit for "high bandwidth" probe transmission-line model, with the response at right.

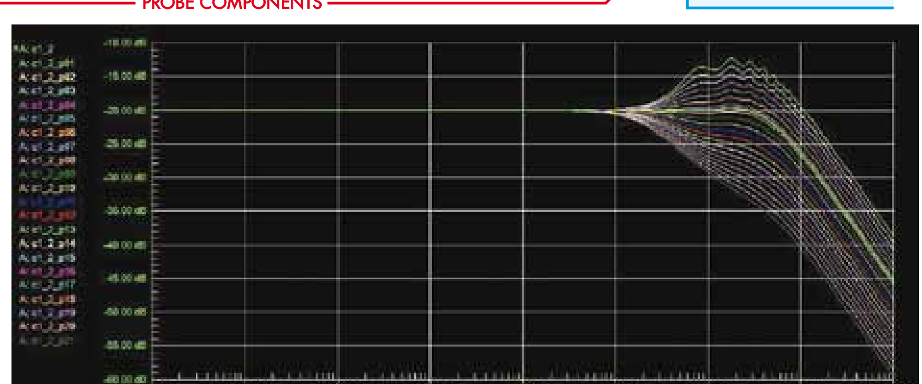

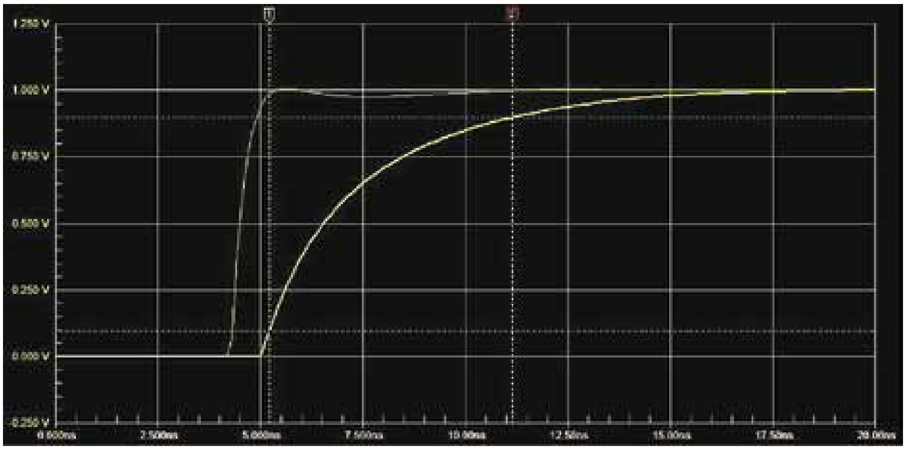

Fig.11: time-domain responses of 60MHz and 600MHz probes.

compensation trim (R_{COMP} in Fig.10) appears to play a significant role. It appears to terminate the lossy transmission line.

For example, if this resistor is shorted, the –3dB bandwidth is reduced to 180MHz and the optimum transmission line resistance is 110Ω/m instead of 115Ω/m. If R_{COMP} is increased above 68Ω to (say) 150Ω, the frequency response shows several dB of peaking at 200MHz.

Interestingly, it makes little difference whether the compensation trimmer and its 68Ω resistor are positioned at the scope end or probe end of the transmission line.

This indicates that the choice of trimmer location is probably a manufacturing decision rather than performance issue.

Rise-time and propagation delay

It is useful to compare the delay (propagation) times of different bandwidth probes. Fig.11 shows the response to a 10V pulse of the 60MHz probe of Fig.9, and the 250MHz probe (with 600MHz bandwidth!) of Fig.10.

The "600MHz" probe (green) has a propagation delay of around 4.2ns while the 60MHz probe (yellow) has around 5.1ns delay.

The propagation delay is the time between the input pulse edge and the start of the pulse edge at the scope end of the cable.

A difference of less than a nanosecond might not seem much, until you're chasing race conditions in logic circuits with mismatched probes.

The rise-time of the scope end waveform is the time taken for the voltage to go from 10% to 90% of the final value. The simulated 60MHz probe shows 5.9ns rise-time; the "600MHz" probe shows 0.7ns rise time.

The effects of faster or slower rise times are in proportion to the nature of the signals you're observing. Nanosecond differences in rise time are irrelevant if you're observing the squarewave response of audio op amps with microsecond rise time but they become vital if you're chasing problems in high-speed digital circuits.

Probe impedance

Does your x10 probe actually have a 10MΩ input impedance? Yes – but only at low frequencies.

Fig.12 shows the input impedance in "dB re 1Ω" of the 60MHz probe of Fig.9. The impedance is 140dB (10MΩ) below 1kHz but the capacitance of the compensation cap determines the impedance at higher frequencies.

It is worth noting that when probing audio circuits at 20kHz, the probe impedance is less than 1MΩ.

At frequencies above the probe's 60MHz bandwidth, the impedance is no longer dominated by the 15pF input capacitance.

It flattens out at 100Ω, dictated by the 50Ω probe tip resistor plus the 50Ω coax impedance.

Probe grounding and ground clips

How "grounded" is the ground clip on your probe?

A typical probe ground wire with alligator clip is around 150mm long. Typical wire inductance is around 1nH/mm, so the ground lead exhibits 150nH of inductance. The probe tip's separation from its ground-lead attachment will add another 50nH or so. This ground inductance was added to the high-bandwidth probe circuit, shown in Fig.13.

The frequency response of this circuit can be compared to the "natural" response of the probe. So our nice, flat 600MHz probe's response has been peaked at 100MHz, with premature rolloff above this.

The transient response isn't pretty either, as seen in Fig.14.

It is worth noting that since most x10 probes have similar input capacitance (10pF to 25pF) and most ground clip leads have a similar length, they will all exhibit peaking around 100MHz, irrespective of probe bandwidth.

For this reason, high-bandwidth probes are generally supplied with a kit of attachments which allow the probe ground to be connected to the circuit via coaxial or other low-inductance paths.

If you're measuring circuit operation above tens of MHz or rise times faster than 50ns, use these fittings!

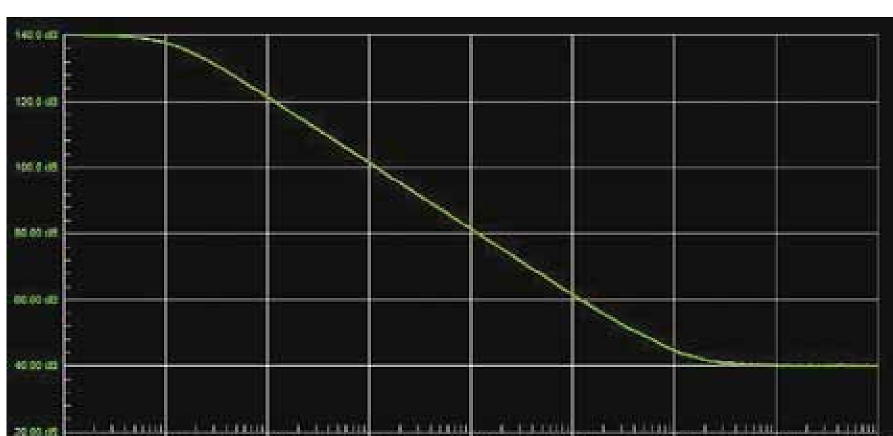

Fig.12: Probe input impedance magnitude.

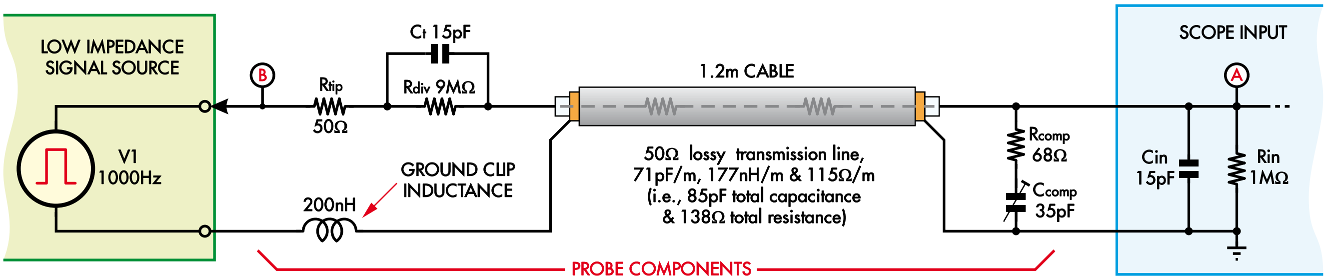

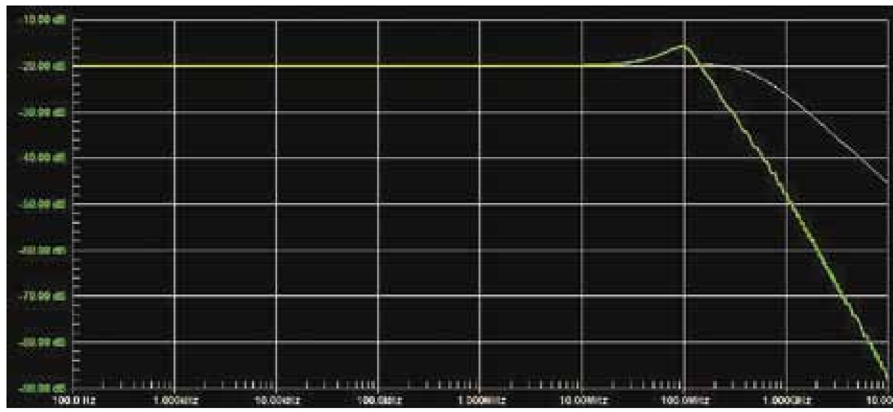

Fig.13: high-bandwidth probe with added ground-clip inductance, with response at right.

Conclusions

The morals of this tale are:

- Trim your probe's compensation capacitor!
- Textbook analyses of probe operation rarely mention transmission-line effects but these are fundamental to the design of a probe.
- There IS a difference between low-bandwidth and high-bandwidth probes. High-bandwidth probes are designed with carefully tailored transmission-line cable and to minimise the effects of end-to-end transmission-line reflections. Much more attention is paid to stray capacitances and build quality.
- A x10 probe will only exhibit 10MΩ impedance at low frequencies. The impedance at higher frequencies is mainly determined by the probe compensating capacitance.
- Use identical probes with equal rise time and bandwidth when inter-channel timing is important (eg, chasing race conditions or clock skew).
- Probe ground-lead inductance can destroy waveform fidelity and bandwidth. Use the kit of adaptors in your probe's pouch to ensure low inductance probe grounds!
- Don't let your probes dangle off the test-bench. Even the good ones break when you tread on them or run your office chair over them!

As a postscript to this article, I received news of the clever fellow who pioneered the use of lossy cable in oscilloscope probes.

It was the invention of John Kobbe, from the halcyon days of Tektronix in the early years. His patent is long expired.

I take my hat off to this gentleman who would have been working without the benefit of PCs and simulator software. SC

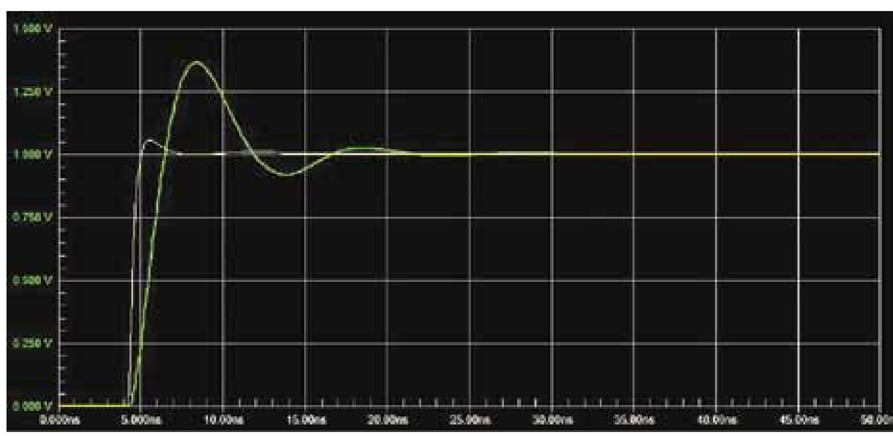

Fig.14: probe waveform with added ground-lead inductance.